

Scalable Latent Tree Model and its Application to Health Analytics

Furong Huang^{*}, Niranjan U.N.^{*}, Ioakeim Perros[†], Robert Chen[†], Jimeng Sun[†], Anima Anandkumar^{*}

Abstract

We present an integrated approach to structure and parameter estimation in latent tree graphical models, where some nodes are hidden. Our overall approach follows a “divide-and-conquer” strategy that learns models over small groups of variables and iteratively merges into a global solution. The structure learning involves combinatorial operations such as minimum spanning tree construction and local recursive grouping; the parameter learning is based on the method of moments and on tensor decompositions. Our method is guaranteed to correctly recover the unknown tree structure and the model parameters with low sample complexity for the class of linear multivariate latent tree models which includes discrete and Gaussian distributions, and Gaussian mixtures. Our bulk asynchronous parallel algorithm is implemented in parallel using the OpenMP framework and scales logarithmically with the number of variables and linearly with dimensionality of each variable. Our experiments confirm a high degree of efficiency and accuracy on large datasets of electronic health records. The proposed algorithm also generates intuitive and clinically meaningful disease hierarchies.

1 Introduction

Latent tree graphical models are a popular class of latent variable models, where a probability distribution involving observed and hidden variables are Markovian on a tree. Due to the fact that structure of (observable and hidden) variable interactions are approximated as a tree, inference on latent trees can be carried out exactly through a simple belief propagation [Pea88]. Therefore, latent tree graphical models present a good trade-off between model accuracy and computational complexity. They are applicable in many domains, where it is natural to expect hierarchical or sequential relationships among the variables (through a hidden-Markov model). For instance, latent tree models have been employed for phylogenetic reconstruction [DEKM99], object recognition [CTW12a, CTW12b] and human pose estimation [WL13]. In this paper, we use latent tree model for discovering a hierarchy among diseases based on comorbidities exhibited in patients’ health records, i.e. co-occurrences of diseases in patients. In particular, two large healthcare datasets of 30K and 1.6M patients are used to build the latent disease trees, where clinically meaningful disease clusters are identified as shown in fig 3 and 4.

The task of learning a latent tree models consists of two parts: learning the tree structure, and learning the parameters of the tree. There exist many challenges which prohibit efficient or guaranteed learning of the latent tree graphical model, which will be addressed in this paper:

1. The location and the number of latent variables are hidden and the marginalized graph over the observable variables no longer conforms to a tree structure.
2. Structure learning algorithms are typically of computational complexity polynomial with p (number of variables) as discussed in [ACH⁺11, CTAW11]. These methods are serial in nature and therefore are not scalable for large p .
3. Parameter estimation in latent tree models is typically carried out through Expectation Maximization (EM) or other local search heuristics [CTAW11]. These methods have no consistency guarantees, suffer from the problem of local optima and are not easily parallelizable.
4. Typically structure learning and parameter estimation are treated sequentially, not together.

^{*}The authors are with Electrical Engineering and Computer Science Dept., University of California, Irvine. Emails: furongh, un.niranjan, a.anandkumar@uci.edu. [†] The authors are with School of Computational Science and Engineering at College of Computing at Georgia Institute of Technology. Emails:perros@gatech.edu, rchen87@gatech.edu, jsun@cc.gatech.edu.

Contributions: In this work, we present an integrated approach to structure and parameter estimation in latent tree models. Our method overcomes all the above shortcomings simultaneously. First, it automatically learns the latent variables and their locations. Second, our method achieves consistent structure estimation with $\log(p)$ computational complexity with enough computational resources via “divide-and-conquer” manner. We also present a rigorous proof on the global consistency of the structure and parameter estimation under the “divide-and-conquer” framework. Our consistency guarantees are applicable to a broad class of linear multivariate latent tree models including discrete distributions, continuous multivariate distributions (e.g. Gaussian), and mixed distributions such as Gaussian mixtures. This model class is much more general than discrete models, prevalent in most of the previous works on latent tree models [MR05, Mos07, ESSW99, AV⁺13]. Third, our algorithm considers the inverse method of moments, and estimates the model parameters via tensor decomposition with low perturbation guarantees. Moreover, we carefully integrate structure learning with parameter estimation, based on tensor spectral decompositions [AGH⁺12]. Finally, our approach has a high degree of parallelism, and is *bulk asynchronous* parallel [GV94].

In addition to the aforementioned technical contributions, we showcase the impact of our work by applying it to two real datasets originating from the healthcare domain. The algorithm was used to discover hidden patterns, or concepts reflecting co-occurrences of particular diagnoses in patients in outpatient and intensive care settings. While such a task is currently done through manual analysis of the data, our method provides an automated method for the discovery of novel clinical concepts from high dimensional, multi-modal data.

Related works: There has been widespread interest in developing distributed learning techniques, e.g. the recent works of [SN10] and [WDK⁺13]. These works consider parameter estimation via likelihood-based optimizations such as Gibbs sampling, while our method involves more challenging tasks where both the structure and the parameters are estimated. Simple methods such as local neighborhood selection through ℓ_1 -regularization [MB06] or local conditional independence testing [ATHW12] can be parallelized, but these methods do not incorporate hidden variables. Finally, note that the latent tree models provide a statistical description, in addition to revealing the hierarchy. In contrast, hierarchical clustering techniques are not based on a statistical model [KBXS12] and cannot provide valuable information such as the level of correlation between observed and hidden variables.

2 Latent Tree Graphical Model Preliminaries

We denote $[n] := \{1, \dots, n\}$. Let $\mathcal{T} := (\mathcal{V}, \mathcal{E})$ denote an undirected tree with vertex set \mathcal{V} and edge set \mathcal{E} . The *neighborhood* of a node v_i , $\text{nb}d(v_i)$, is the set of nodes to which v_i is directly connected on the tree. Leaves which have a common neighboring node are known as *siblings*, and the common node is referred to as their *parent*. Let N denote the number of samples. An example of latent tree is depicted in Figure 1(a).

There are two types of variables on the nodes, namely, the observable variables, denoted by $\mathcal{X} := \{x_1, \dots, x_p\}$ ($p := |\mathcal{X}|$), and hidden variables, denoted by $\mathcal{H} := \{h_1, \dots, h_m\}$ ($m := |\mathcal{H}|$). Let $\mathcal{Y} := \mathcal{X} \cup \mathcal{H}$ denote the complete set of variables and let y_i denote the random variable at node $v_i \in \mathcal{V}$, and similarly let y_A denote the set of random variables in set A .

A graphical model is defined as follows: given the neighborhood $\text{nb}d(v_i)$ of any node $v_i \in \mathcal{V}$, the variable y_i is conditionally independent of the rest of the variables in \mathcal{V} , i.e., $y_i \perp\!\!\!\perp y_j | y_{\text{nb}d(v_i)}$, $\forall v_j \in \mathcal{V} \setminus \{v_i \cup \text{nb}d(v_i)\}$.

Linear Models: We consider the class of linear latent tree models. The observed variables x_i are random vectors of length d_i , i.e., $x_i \in \mathbb{R}^{d_i}$, $\forall i \in [p]$ while the latent nodes are k -state categorical variables, i.e., $h_i \in \{e_1, \dots, e_k\}$, where $e_j \in \mathbb{R}^k$ is the j^{th} standard basis vector. Although d_i can vary across variables, we use d for notation simplicity. In other words, for notation simplicity, $x_i \in \mathbb{R}^d$, $\forall i \in [p]$ is equivalent to $x_i \in \mathbb{R}^{d_i}$, $\forall i \in [p]$. For any variable y_i with neighboring hidden variable h_j , we assume a linear relationship:

$$\mathbb{E}[y_i | h_j] = A_{y_i | h_j} h_j, \quad (1)$$

where transition matrix $A_{y_i | h_j} \in \mathbb{R}^{d \times k}$ is assumed to have full column rank, $\forall y_i, h_j \in \mathcal{Y}$. This implies that $k \leq d$, which is natural if we want to enforce a parsimonious model for fitting the observed data.

For a pair of (observed or hidden) variables y_a and y_b , consider the *pairwise correlation matrix* $\mathbb{E}[y_a y_b^\top]$ where the expectation is over samples. Since our model assumes that two observable variables interact through at least a hidden variable, we have

$$\mathbb{E}[y_a y_b^\top] := \sum_{e_i} \mathbb{E}[h_j = e_i] A_{y_a | h_j = e_i} A_{y_b | h_j = e_i}^\top \quad (2)$$

We see that $\mathbb{E}[y_a y_b^\top]$ is of rank k since $A_{y_a | h_j = e_i}$ or $A_{y_b | h_j = e_i}$ is of rank k .

3 Overview of Approach

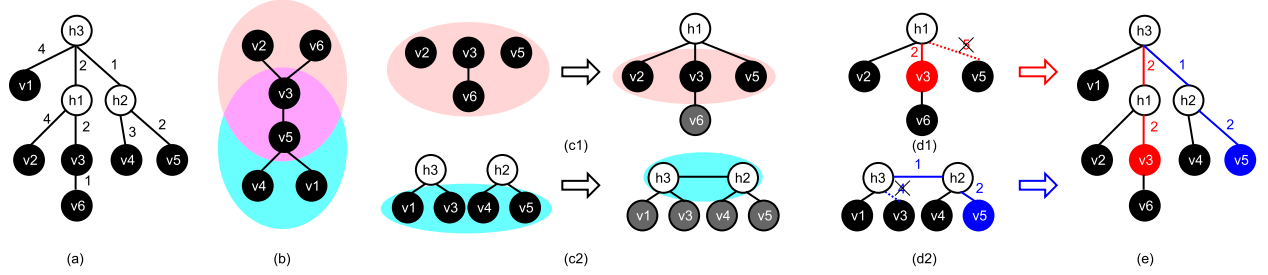


Figure 1: (a) Ground truth latent tree to be estimated, numbers on edges are *multivariate information distances*. (b) MST constructed using the *multivariate information distances*. v_3 and v_5 are internal nodes (leaders). Note that *multivariate information distances* are additive on latent tree, not on MST. (c1) LCR on $\text{nbd}[v_3, \text{MST}]$ to get local structure \mathcal{N}_3 . Pink shadow denotes the active set. Local parameter estimation is carried out over triplets with joint node, such as (v_2, v_3, v_5) with joint node h_1 . (c2) LCR on $\text{nbd}[v_5, \text{MST}]$ to get local structure \mathcal{N}_5 . Cyan shadow denotes the active set. (d1)(d2) Merging local sub-trees. $\text{Path}(v_3, v_5; \mathcal{N}_3)$ and $\text{path}(v_3, v_5; \mathcal{N}_5)$ conflict. (e) Final recovery.

The overall approach is depicted in Figure 1, where (a) and (b) show the data preprocessing step, (c) - (e) illustrate the divide-and-conquer step for structure and parameter learning.

More specifically, we start with the parallel computation of pairwise *multivariate information distances*. Information distance roughly measures the extent of correlation between different pairs of observed variables and requires SVD computations in step (a). Then in step (b) a Minimum Spanning Tree (MST) is constructed over observable variables in parallel [BC06] using the *multivariate information distance*. The local groups are also obtained through MST so that they are available for the structure and parameter learning step that follows.

The structure and parameter learning is done jointly through a divide-and-conquer strategy. Step-(c) illustrates the divide step (or local learning), where local structure and parameter estimation is performed. It also performs the local merge to obtain group level structure and parameter estimates. After the local structure and parameter learning is finished within the groups, we perform merge operations among groups, again guided by the Minimum Spanning Tree structure. For the structure estimation it consists of a union operation of sub-trees; for the parameter estimation, it consists of linear algebraic operations. Since our method is unsupervised, an alignment procedure of the hidden states is carried out which finalizes the global estimates of the tree structure and the parameters.

4 Structure Learning

Structure learning in graphical models involves finding the underlying Markov graph, given the observed samples. For latent tree models, structure can be estimated via distance based methods. This involves computing certain *information distances* between any pair of observed variables, and then finding a tree which fits the computed distances.

Multivariate information distances: We propose an additive distance for multivariate linear latent tree models. For a pair of (observed or hidden) variables y_a and y_b , consider the pairwise correlation matrix $\mathbb{E}[y_a y_b^\top]$ (the expectation is over samples). Note that its rank is k , dimension of the hidden variables.

Definition 4.1. The multivariate information distance between nodes i and j is defined as

$$\text{dist}(v_a, v_b) := -\log \frac{\prod_{i=1}^k \sigma_i(\mathbb{E}(y_a y_b^\top))}{\sqrt{\det(\mathbb{E}(y_a y_a^\top)) \det(\mathbb{E}(y_b y_b^\top))}} \quad (3)$$

where $\{\sigma_1(\cdot), \dots, \sigma_k(\cdot)\}$ are the top k singular values.

Note that definition 4.1 suggests that this multivariate information distance allows heterogeneous settings where the dimensions of y_a and y_b are different (and $\geq k$).

For latent tree models, we can find information distances which are provably *additive* on the underlying tree in expectation, i.e. the expected distance between any two nodes in the tree is the sum of distances along the path between them.

Lemma 4.2. The multivariate information distance is additive on the tree \mathcal{T} , i.e., $\text{dist}(v_a, v_c) = \text{dist}(v_a, v_b) + \text{dist}(v_b, v_c)$, where v_b is a node in the path from v_a to v_c and $v_a, v_b, v_c \in \mathcal{V}$.

Refer to Appendix A for proof. The empirical distances can be computed via rank- k SVD of the empirical pairwise moment matrix $\hat{\mathbb{E}}[y_a y_b^\top]$. Note that the distances for all the pairs can be computed in parallel.

Formation of local groups via MST: Once the empirical distances are computed, we construct a Minimum Spanning Tree (MST), based on those distances. Note that the MST can be computed efficiently in parallel [VHPN09, Mic12]. We now form groups of observed variables over which we carry out learning independently, without any coordination. These groups are obtained by the (closed) neighborhoods in the MST, i.e. an internal node and its one-hop neighbors form a group. The corresponding internal node is referred to as the *group leader*. See Figure 1(b).

Local recursive grouping (LRG): Once the groups are constructed via neighborhoods of MST, we construct a sub-tree with hidden variables in each group (in parallel) using the recursive grouping introduced in [CTAW11]. The recursive grouping uses the multivariate information distances and decides the locations and numbers of hidden nodes. It proceeds by deciding which nodes are siblings, which proceeds as follows: consider two observed nodes v_i, v_j which are siblings on the tree with a common parent v_l , and consider any other observed node v_a . From additivity of the (expected) information distances, we have $\text{dist}(v_i, v_a) = \text{dist}(v_i, v_l) + \text{dist}(v_l, v_a)$ and similarly for $\text{dist}(v_j, v_a)$. Thus, we have $\Phi(v_i, v_j; v_a) := \text{dist}(v_i, v_a) - \text{dist}(v_j, v_a) = \text{dist}(v_i, v_l) - \text{dist}(v_j, v_l)$, which is independent of node v_a . Thus, comparing the quantity $\Phi(v_i, v_j; v_a)$ for different nodes v_a allows us to conclude that v_i and v_j are siblings. Once the siblings are inferred, the hidden nodes are introduced, and the same procedure repeats to construct the higher layers. Note that whenever we introduce a new hidden node h_{new} as a parent, we need to estimate multivariate information distance between h_{new} and nodes in active set Ω . This is discussed in [CTAW11] with details.

We will describe the LRG in details with integrated parameters estimation in Procedure 1 in Section 6. In the end, we obtain a sub-tree over the local group of variables. After this *local recursive grouping test*, we store the neighborhood relationship for the leader v_i using an adjacency list \mathcal{N}_i . We call the resultant local structure as *latent sub-tree*.

5 Parameter Estimation

Along with the structure learning, we adopt a moment-based spectral learning technique for parameter estimation. This is a guaranteed and fast approach to recover parameters via moment matching for third order moments of the observed data. In contrast, traditional approaches such as Expectation Maximization (EM) suffer from spurious local optima and cannot provably recover the parameters.

A latent tree with three leaves: We first consider an example of three observable leaves x_1, x_2, x_3 (i.e., a triplet) with a common hidden parent h . We then clarify how this can be generalized to learn the parameters of the latent tree model. Let \otimes denote for the tensor product. For example, if $x_1, x_2, x_3 \in \mathbb{R}^d$, we have $x_1 \otimes x_2 \otimes x_3 \in \mathbb{R}^{d \times d \times d}$. **Property 5.1** (tensor decomposition for triplets). For a linear latent tree model with three observed nodes v_1, v_2, v_3 with joint hidden node h , we have

$$\mathbb{E}(x_1 \otimes x_2 \otimes x_3) = \sum_{r=1}^k \mathbb{P}[h = e_r] A_{x_1|h}^r \otimes A_{x_2|h}^r \otimes A_{x_3|h}^r, \quad (4)$$

where $A_{x_i|h}^r = \mathbb{E}(x_i|h = e_r)$, i.e., r^{th} column of the transition matrices from h to x_i . The tensor decomposition method of [AGH⁺12] provably recovers the parameters $A_{x_i|h}$, $\forall i \in [3]$, and $\mathbb{P}[h]$.

Tensor decomposition for learning latent tree models: We employ the above approach for learning latent tree model parameters as follows: for every triplet of variables y_a , y_b , and y_c (hidden or observed), we consider the hidden variable h_i which is the joining point of y_a , y_b and y_c on the tree. They form a *triplet* model, for which we employ the tensor decomposition procedure. However, it is wasteful to do it over all the triplets in the latent tree.

In the next section, we demonstrate how we efficiently estimate the parameters as we learn the structure, and minimize the tensor decompositions required for estimation. Issues such as alignment of hidden labels across different decompositions will also be addressed.

6 Integrated Struct. and Param. Estimation

So far, we described high-level procedures of structure estimation through local recursive grouping (LRG) and parameter estimation through tensor decomposition over triplets of variables, respectively. We now describe an integrated and efficient approach which brings all these ingredients together. In addition, we provide merging steps to obtain a global model, using the sub-trees and parameters learnt over local groups.

6.1 Local recursive grouping with tensor decomp.

Next we present an integrated procedure where the parameter estimation goes hand-in-hand with structure estimation. Intuitively, we find efficient groups of triplets to carry out tensor decomposition simultaneously, as we estimate the structure through recursive grouping. In recursive grouping, pairs of nodes are recursively grouped as siblings or as parent-child. As this process continues, we carry out tensor decompositions whenever there are siblings present as triplets. If there are only a pair of siblings, we find an observed node with closest distance to the pair. Once the tensor decompositions are carried out on the observed nodes, we proceed to structure and parameter estimation of the added hidden variables. The samples of the hidden variables can be obtained via the posterior distribution, which is learnt earlier through tensor decomposition. This allows us to predict information distances and third order moments among the hidden variables as process continues. The full algorithm is given in Procedure 1.

Procedure 1 LRG with Parameter Estimation

Input: for each $v_i \in \mathcal{X}_{\text{int}}$, active set $\Omega := \text{nbd}[v_i; \text{MST}]$.

Output: for each $v_i \in \mathcal{X}_{\text{int}}$, local sub-tree adjacency matrix \mathcal{N}_i , and $\mathbb{E}[y_a|y_b]$ for all $(v_a, v_b) \in \mathcal{N}_i$.

- 1: Active set $\Omega \leftarrow \text{nbd}[v_i; \text{MST}]$
 - 2: **while** $|\Omega| > 2$ **do**
 - 3: **for all** $v_a, v_b \in \Omega$ **do**
 - 4: **if** $\Phi(v_a, v_b; v_c) = \text{dist}(v_a, v_b)$, $\forall v_c \in \Omega \setminus \{v_a, v_b\}$ **then**
 - 5: v_a is a leaf node and v_b is its parent,
 - 6: Eliminate v_a from Ω .
 - 7: **if** $-\text{dist}(v_a, v_b) < \Phi(v_a, v_b; v_c) = \Phi(v_a, v_b; v'_c) < \text{dist}(v_a, v_b)$, $\forall v_c, v'_c \in \Omega \setminus \{v_a, v_b\}$ **then**
 - 8: v_a and v_b are siblings, eliminate v_a and v_b from Ω , add h_{new} to Ω .
 - 9: Introduce new hidden node h_{new} as parent of v_a and v_b .
 - 10: **if** more than 3 siblings under h_{new} **then**
 - 11: find v_c in siblings,
 - 12: **else**
 - 13: find $v_c = \arg \min_{v_c \in \Omega} \text{dist}(v_a, v_c)$.
 - 14: Estimate empirical third order moments $\hat{\mathbb{E}}(y_a \otimes y_b \otimes y_c)$
 - 15: Decompose $\hat{\mathbb{E}}(y_a \otimes y_b \otimes y_c)$ to get $\text{Pr}[h_{\text{new}}]$ and $\mathbb{E}(y_r|h_{\text{new}})$, $\forall r = \{a, b, c\}$.
-

The divide-and-conquer local spectral parameter estimation is superior compared to popular EM-based method [CTAW11], which is slow and prone to local optima. More importantly, EM can only be applied on a stable structure since it is a

global update procedure. Our proposed spectral learning method, in contrast, is applied locally over small groups of variables, and is a guaranteed learning with sufficient number of samples [AGH⁺12]. Moreover, since we integrate structure and parameter learning, we avoid recomputing the same quantities, e.g. SVD computations are required both for structure estimation (for computing distances) and parameter estimation (for whitening the tensor). Combining these operations results in huge computational savings (see Section 7 for the exact computational complexity of our method).

Procedure 2 Merging and Alignment Correction (MAC)

Input: Latent sub-trees \mathcal{N}_i for all internal nodes i .

Output: Global latent tree T structure and parameters.

- 1: **for** \mathcal{N}_i and \mathcal{N}_j in all the sub-trees **do**
 - 2: **if** there are common nodes between \mathcal{N}_i and \mathcal{N}_j **then**
 - 3: Find the shortest path $\text{path}(v_i, v_j; \mathcal{N}_i)$ between v_i and v_j on \mathcal{N}_i and $\text{path}(v_i, v_j; \mathcal{N}_j)$ in \mathcal{N}_j ;
 - 4: Union the only conflicting path $\text{path}(v_i, v_j; \mathcal{N}_i)$ and $\text{path}(v_i, v_j; \mathcal{N}_j)$ according to equation (7) ;
 - 5: Attach other nodes in \mathcal{N}_i and \mathcal{N}_j to the union path;
 - 6: Perform alignment correction as described in Procedure 3.
-

6.2 Merging and Alignment Correction

We have so far learnt sub-trees and parameters over local groups of variables, where the groups are determined by the neighborhoods of the MST. The challenge now is to combine them to obtain a globally consistent estimate. There are non-trivial obstacles to achieving this: first, the constructed local sub-trees span overlapping groups of observed nodes, and possess conflicting paths. Second, local parameters need to be re-aligned as we merge the subtrees to obtain globally consistent estimates due to the nature of unsupervised learning. To be precise, different tensor decompositions lead to permutation of the hidden labels (i.e. columns of the transition matrices) across triplets. Thus, we need to find the permutation matrix correcting the alignment of hidden states of the transition matrices, so as to guarantee global consistency.

Structure Union: We now describe the procedure to merge the local structures. We merge them in pairs to obtain the final global latent tree. Recall that \mathcal{N}_i denotes a sub-tree constructed locally over a group, whose leader is node v_i . Consider a pair of subtrees \mathcal{N}_i and \mathcal{N}_j , whose group leaders v_i and v_j are neighbors on the MST. Since v_i and v_j are neighbors, both the sub-trees contain them, and have different paths between them (with hidden variables added). Moreover, note that this is the only conflicting path in the two subtrees. We now describe how we can resolve this: in \mathcal{N}_i , let h_1^i be the neighboring hidden node for v_i and h_2^i be the neighbor of v_j . There could be more hidden nodes between h_1^i and h_2^i . Similarly, in \mathcal{N}_j , let h_1^j and h_2^j be the corresponding nodes in \mathcal{N}_j . The shortest path between v_i and v_j in the two sub-trees are given as follows:

$$\text{path}(v_i, v_j; \mathcal{N}_i) := [v_i - h_1^i - \dots - h_2^i - v_j] \quad (5)$$

$$\text{path}(v_i, v_j; \mathcal{N}_j) := [v_i - h_1^j - \dots - h_2^j - v_j] \quad (6)$$

Then the union path is formed as follows:

$$\begin{aligned} &\text{merge}(\text{path}(v_i, v_j; \mathcal{N}_i), \text{path}(v_i, v_j; \mathcal{N}_j)) \\ &:= [v_i - h_1^i - \dots - h_2^i - h_1^j \dots h_2^j - v_j] \end{aligned} \quad (7)$$

In other words, we retain the immediate hidden neighbor of each group leader, and break the paths on the other end. For example in Figure 1(d1,d2), we have the path $v_3 - h_1 - v_5$ in \mathcal{N}_3 and path $v_3 - h_3 - h_2 - v_5$ in \mathcal{N}_5 . The resulting path is $v_3 - h_1 - h_3 - h_2 - v_5$, as see in Figure 1(e). After the union of the conflicting paths, the other nodes are attached to the resultant latent tree. We present the pseudo code in Procedure 2 in Appendix E.

Procedure 3 Parameter Alignment Correction

(\mathbb{G}_r denotes reference group, \mathbb{G}_o denotes the list of other groups, each group has a reference node denoted as \mathcal{R}_l , and the reference node in \mathbb{G}_r is \mathcal{R}_g . The details on alignment at line 8 is in Appendix E.)

Input: Triplets and unaligned parameters estimated for these triplets, denoted as $\text{Trip}(y_i, y_j, y_k)$.

Output: Aligned parameters for the entire latent tree T .

- 1: Select \mathbb{G}_r which has *sufficient children*;
 - 2: Select refer node \mathcal{R}_g in \mathbb{G}_r ;
 - 3: **for all** a, b in \mathbb{G}_r **do**
 - 4: Align $\text{Trip}_{\text{in}}(y_a, y_b, \mathcal{R}_g)$;
 - 5: **for all** i_g in \mathbb{G}_o **do**
 - 6: Select refer node \mathcal{R}_l in $\mathbb{G}_o[i_g]$;
 - 7: Align $\text{Trip}_{\text{out}}(\mathcal{R}_g, y_a, \mathcal{R}_l)$ and $\text{Trip}_{\text{out}}(\mathcal{R}_l, y_i, \mathcal{R}_g)$;
 - 8: **for all** i, j in $\mathbb{G}_o[i_g]$ **do**
 - 9: Align $\text{Trip}(y_i, y_j, \mathcal{R}_l)$;
-

Parameter Alignment Correction: As mentioned before, our parameter estimation is unsupervised, and therefore, columns of the estimated transition matrices may be permuted for different triplets over which tensor decomposition is carried out. Note that the parameter estimation within the triplet is automatically acquired through the tensor decomposition technique, so that the alignment issue only arises across triplets. We refer to this as the alignment issue and it is required at various levels.

There are two types of triplets, namely, *in-group* and *out-group* triplets. A triplet of nodes $\text{Trip}(y_i, y_j, y_l)$ is said to be *in-group* (denoted by $\text{Trip}_{\text{in}}(y_i, y_j, y_l)$) if its containing nodes share a joint node h_k and there are no other hidden nodes in $\text{path}(y_i, h_k)$, $\text{path}(y_j, h_k)$ or $\text{path}(y_l, h_k)$. Otherwise, this triplet is *out-group* denoted by $\text{Trip}_{\text{out}}(y_i, y_j, y_l)$. We define a group as *sufficient children* group if it contains at least three *in-group* nodes.

Designing an *in-group* alignment correction with *sufficient children* is relatively simple: we achieve this by including a local reference node for all the *in-group* triplets. Thus, all the triplets are aligned with the reference node. The alignment correction is more challenging if lacking *sufficient children*. We propose *out-group* alignment to solve this problem. We first assign one group as a *reference group*, and the *local reference node* in that *reference group* becomes the *global reference node*. In this way, we align all recovered transition matrices in the same order of hidden states as in the reference node. Overall, we merge the local structures and align the parameters from LRG local sub-trees using Procedure 2 and 3.

7 Theoretical Gaurantees

Correctness of Proposed Parallel Algorithm: We now provide the main result of this paper on global consistency for our method, despite the high degree of parallelism.

Theorem 7.1. *Given samples from an identifiable latent tree model, the proposed method consistently recovers the structure with $O(\log p)$ sample complexity and parameters with $O(\text{poly } p)$ sample complexity.*

The proof sketch is in Appendix C.

Computational Complexity: We recall some notations here: d is the observable node dimension, k is the hidden node dimension ($k \ll d$), N is the number of samples, p is the number of observable nodes, and z is the number of non-zero elements in each sample.

Let Γ denote the maximum size of the groups, over which we operate the local recursive grouping procedure. Thus, Γ affects the degree of parallelism for our method. Recall that it is given by the neighborhoods on MST, i.e., $\Gamma := \max_i |\text{nbr}[i; \text{MST}]|$. Below, we provide a bound on Γ .

Lemma 7.2. *The maximum size of neighborhoods on MST, denoted as Γ , satisfies*

$$\Gamma \leq \Delta^{1 + \frac{zd}{\Gamma d}}, \quad (8)$$

where $\delta := \max_i \{\min_j \{\text{path}(v_i, v_j; \mathcal{T})\}\}$ is the effective depth, Δ is the maximum degree of \mathcal{T} , and the u_d and l_d are the upper and lower bound of information distances between neighbors on \mathcal{T} .

Thus, we see that for many natural cases, where the degree and the depth in the latent tree are bounded (e.g. the hidden Markov model), and the parameters are mostly homogeneous (i.e., u_d/l_d is small), the group sizes are bounded, leading to a high degree of parallelism.

We summarize the computational complexity in Table 1. Details can be found in Appendix F.

Algorithm Steps	Time per worker	Degree of parallelism
Distance Est.	$O(Nz + d + k^3)$	$O(p^2)$
MST	$O(\log p)$	$O(p^2)$
LRG	$O(\Gamma^3)$	$O(p/\Gamma)$
Tensor Decomp.	$O(\Gamma k^3 + \Gamma dk^2)$	$O(p/\Gamma)$
Merging step	$O(dk^2)$	$O(p/\Gamma)$

Table 1: Worst-case computational complexity of our algorithm. The total complexity is the product of the time per work and degree of parallelism.

8 Experiments

Setup Experiments are conducted on a server running the Red Hat Enterprise 6.6 with 64 AMD Opteron processors and 265 GBRAM. The program is written in C++, coupled with the multi-threading capabilities of the OpenMP environment [DM98] (version 1.8.1). We use the Eigen toolkit¹ where BLAS operations are incorporated. For SVDs of large matrices, we use randomized projection methods [GM13a] as described in Appendix H.

Healthcare data analysis The goal of our analysis is to discover a disease hierarchy based on their co-occurring relationships in the patient records. In general, longitudinal patient records store the diagnosed diseases on patients over time, where the diseases are encoded with International Classification of Diseases (ICD) code.

Data description We used two large patient datasets of different sizes with respect to the number of samples, variables and dimensionality.

(1) *MIMIC2*: The MIMIC2 dataset record disease history of 29,862 patients where a overall of 314,647 diagnostic events over time representing 5675 diseases are logged. We consider patients as samples and groups of diseases as variables. We analyze and compare the results by varying the group size (therefore varying d and p).

(2) *CMS*: The CMS dataset includes 1.6 million patients, for whom 15.8 million medical encounter events are logged. Across all events, 11,434 distinct diseases (represented by ICD codes) are logged. We consider patients as samples and groups of diseases as variables. We consider specific diseases within each group as dimensions. We analyze and compare the results by varying the group size (therefore varying d and p). While the MIMIC2 dataset and CMS dataset both contain logged diagnostic events, the larger volume of data in CMS provides an opportunity for testing the algorithm’s scalability. We qualitatively evaluate biological implications on MIMIC2 and quantitatively evaluate algorithm performance and scalability on CMS.

To learn the disease hierarchy from data, we also leverage some existing domain knowledge about diseases. In particular, we use an existing mapping between ICD codes and higher-level Phenome-wide Association Study (PheWAS) codes [DRB⁺10]. We use (about 200) PheWAS codes as observed nodes and the observed node dimension is set to be binary ($d = 2$) or the maximum number of ICD codes within a pheWAS code ($d = 31$). The goal is to learn the latent nodes and the disease hierarchy and associated parameters from data.

8.1 Experimental results - Validation

We conduct both quantitative and qualitative validation of the resulting disease hierarchy.

¹http://eigen.tuxfamily.org/index.php?title=Main_Page




Figure 2: (a) CMS dataset sub-sampling w.r.t. varying number of samples. (b) MIMIC2 dataset sub-sampling w.r.t. varying number of observed nodes. Each one of the observed nodes is binary ($d = 2$). (c) MIMIC2 dataset: Scaling w.r.t. varying computational power, establishing the scalability of our method even in the large p regime. The number of observed nodes is 1083 and each one of them is binary ($p = 1083, d = 2$).

Quantitative Analysis We first compare our resulting hierarchy with a ground truth tree based on medical knowledge². The standard Robinson Foulds (RF) metric [RF81](between our estimated latent tree and the ground truth tree) is computed to evaluate the structure recovery in Table 2. The smaller the metric is, the better the recovered tree is. We also compare our results with a baseline: the agglomerative clustering. The proposed method are slightly better than the baseline and the advantage is increased with more nodes. However, the proposed method provides an efficient probabilistic graphical model that can support general inference which is beyond the baseline.

Data	p	RF(agglo.)	RF(proposed)
MIMIC2	163	0.0061	0.0061
CMS	168	0.0060	0.0059
MIMIC2	952	0.0060	0.0011

Table 2: Robinson Foulds (RF) metric compared with the “ground-truth” tree for both MIMIC2 and CMS dataset. Our proposed results are better as we increase the number of nodes.

Qualitative analysis The qualitative analysis is done by a senior MD-PhD student in our team.

(a) Case $d=2$: Here we report the results from the 2-dimensional case (i.e., observed variable is binary). In figure 3, we show a portion of the learned tree using the MIMIC2 healthcare data. The yellow nodes are latent nodes from the learned subtrees while the blue nodes represent observed nodes(diagnosis codes) in the original dataset. Diagnoses that are similar were generally grouped together. For example, many neoplastic diseases were grouped under the same latent node (node 1135). While some dissimilar diseases were grouped together, there usually exists a known or plausible association of the diseases in the clinical setting. For example, in figure 3, clotting-related diseases and altered mental status were grouped under the same latent node as several neoplasms. This may reflect the fact that altered mental status and clotting conditions such as thrombophlebitis can occur as complications of neoplastic diseases [FMVB03]. The association of malignant neoplasms of prostate and colon polyps, two common cancers in males, is captured under latent node 1136 [G⁺14].

(b) Case $d=31$: We also learn a tree from the MIMIC2 dataset, in which we grouped diseases into 163 pheWAS codes and up to 31 dimensions per variable. Figure 4 shows a portion of the learned tree of four subtrees which all reflect similar diseases relating to trauma. A majority of the learned subtrees reflected clinically meaningful concepts, in that related and commonly co-occurring diseases tended to group together in the same subtrees or in nearby subtrees. We also learn the disease tree from the larger CMS dataset, in which we group diseases into 168 variables and up to 31 dimensions per variable. Similar to the case from the MIMIC2 dataset, a majority of learned subtrees reflected clinically meaningful concepts.

²The ground truth tree is the PheWAS hierarchy provided in the clinical study [DRB⁺10]

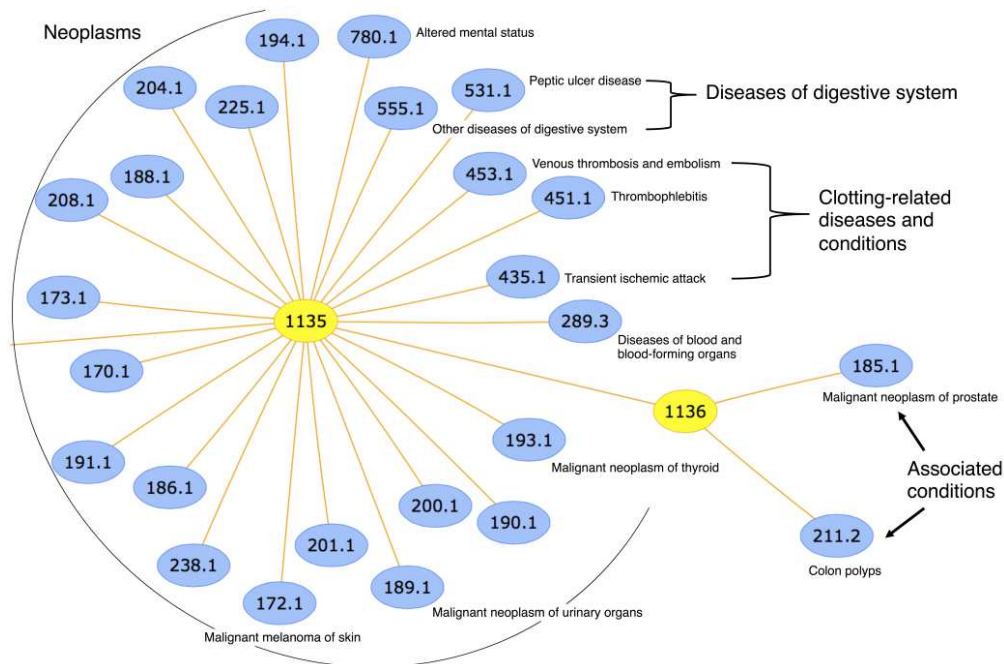


Figure 3: An example of two subtrees which represent groups of similar diseases which may commonly co-occur. Nodes colored yellow are latent nodes from learned subtrees.

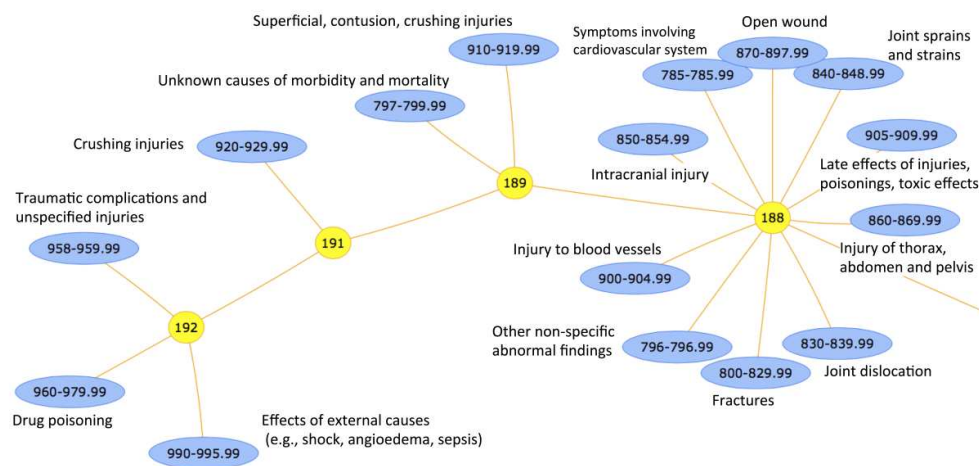


Figure 4: An example of four subtrees which represent groups of similar diseases which may commonly co-occur. Most variables in this subtree are related to trauma.

For both the MIMIC2 and CMS datasets, we performed a qualitative comparison of the resulting trees while varying the hidden dimension k for the algorithm. The resulting trees for different values of k did not exhibit significant differences. This implies that our algorithm is robust with different choices of hidden dimensions. The estimated model parameters are also robust for different values of k based on the results.

Scalability Our algorithm is scalable w.r.t. varying characteristics of the input data. First, it can handle a large number of patients efficiently, as shown in Figure 2(a). It has also a linear scaling behavior as we vary the number observed nodes, as shown in Figure 2(b). Furthermore, even in cases where the number of observed variables is large, our method maintains an almost linear scale-up as we vary the computational power available, as shown in Figure 2(c). As such, by providing the respective resources, our algorithm is practical under any variation of the input data characteristics.

9 Conclusion

We present an integrated approach to structure and parameter estimation in latent tree models. Our method overcomes challenges such as uncertainty of location and number of hidden variables, problem of local optima with no consistency guarantees, difficulty in scalability with respect to number of variables. The proposed algorithm is ideal for parallel computing and highly scalable. We successfully applied the algorithm to a real application for disease hierarchy discovery using large patient data for 1.6m patients.

Acknowledgement

The first author is supported by NSF BIGDATA IIS-1251267, the second author is supported in part by UCI graduate fellowship and NSF Award CCF-1219234, and the last author is supported in part by Microsoft Faculty Fellowship, NSF Career award CCF-1254106, NSF Award CCF-1219234, and ARO YIP Award W911NF-13-1-0084.

References

- [ACH⁺11] Animashree Anandkumar, Kamalika Chaudhuri, Daniel Hsu, Sham M Kakade, Le Song, and Tong Zhang. Spectral methods for learning multivariate latent tree structure. *arXiv preprint arXiv:1107.1283*, 2011.
- [AFH⁺12] Animashree Anandkumar, Dean P Foster, Daniel Hsu, Sham M Kakade, and Yi-Kai Liu. Two svds suffice: Spectral decompositions for probabilistic topic modeling and latent dirichlet allocation. *CoRR*, abs/1204.6703, 1, 2012.
- [AGH⁺12] Anima Anandkumar, Rong Ge, Daniel Hsu, Sham M Kakade, and Matus Telgarsky. Tensor decompositions for learning latent variable models. *arXiv preprint arXiv:1210.7559*, 2012.
- [ATHW12] A. Anandkumar, V. Y. F. Tan, F. Huang, and A. S. Willsky. High-dimensional structure learning of Ising models: local separation criterion. *The Annals of Statistics*, 40(3):1346–1375, 2012.
- [AV⁺13] Animashree Anandkumar, Ragupathyraj Valluvan, et al. Learning loopy graphical models with latent variables: Efficient methods and guarantees. *The Annals of Statistics*, 41(2):401–435, 2013.
- [BC06] David A Bader and Guojing Cong. Fast shared-memory algorithms for computing the minimum spanning forest of sparse graphs. *Journal of Parallel and Distributed Computing*, 66(11):1366–1378, 2006.
- [CTAW11] Myung Jin Choi, Vincent YF Tan, Animashree Anandkumar, and Alan S Willsky. Learning latent tree graphical models. *The Journal of Machine Learning Research*, 12:1771–1812, 2011.
- [CTW12a] M.J. Choi, A. Torralba, and A.S. Willsky. Context models and out-of-context objects. *Pattern Recognition Letters*, 2012.

- [CTW12b] Myung Jin Choi, Antonio Torralba, and Alan S Willsky. Context models and out-of-context objects. *Pattern Recognition Letters*, 33(7):853–862, 2012.
- [CW13] Kenneth L Clarkson and David P Woodruff. Low rank approximation and regression in input sparsity time. In *Proceedings of the 45th annual ACM symposium on Symposium on theory of computing*, pages 81–90. ACM, 2013.
- [DEKM99] R. Durbin, S. R. Eddy, A. Krogh, and G. Mitchison. *Biological Sequence Analysis: Probabilistic Models of Proteins and Nucleic Acids*. Cambridge Univ. Press, 1999.
- [DM98] Leonardo Dagum and Ramesh Menon. Openmp: an industry standard api for shared-memory programming. *Computational Science & Engineering, IEEE*, 5(1):46–55, 1998.
- [DRB⁺10] JC Denny, MD Ritchie, MA Basford, JM Pulley, L Bastarache, K Brown-Gentry, D Wang, DR Masys, Roden DM, and DC Crawford. Phewas: demonstrating the feasibility of a phenome-wide scan to discover genedisease associations. *Bioinformatics*, 26(9):1205–1210, 2010.
- [ESSW99] Peter L Erdos, Michael A Steel, László A Székely, and Tandy J Warnow. A few logs suffice to build (almost) all trees (i). *Random Structures and Algorithms*, 14(2):153–184, 1999.
- [FMVB03] A Falanga, M Marchetti, A Vignoli, and D Balducci. Clotting mechanisms and cancer: implications in thrombus formation and tumor progression. *Clinical advances in hematology & oncology: H&O*, 1(11):673–678, 2003.
- [G⁺14] US Cancer Statistics Working Group et al. United states cancer statistics: 1999–2010 incidence and mortality web-based report. *Atlanta (GA): Department of Health and Human Services, Centers for Disease Control and Prevention, and National Cancer Institute*, 2014.
- [GM13a] Alex Gittens and Michael W Mahoney. Revisiting the nystrom method for improved large-scale machine learning. *arXiv preprint arXiv:1303.1849*, 2013.
- [GM13b] Alex Gittens and Michael W. Mahoney. Revisiting the nystrom method for improved large-scale machine learning. *CoRR*, abs/1303.1849, 2013.
- [GV94] Alexandros V Gerbessiotis and Leslie G Valiant. Direct bulk-synchronous parallel algorithms. *Journal of parallel and distributed computing*, 22(2):251–267, 1994.
- [HNH⁺13] Furong Huang, Niranjan U. N, Mohammad Umar Hakeem, Prateek Verma, and Animashree Anandkumar. Fast detection of overlapping communities via online tensor methods on gpus. *CoRR*, abs/1309.0787, 2013.
- [KBXS12] Akshay Krishnamurthy, Sivaraman Balakrishnan, Min Xu, and Aarti Singh. Efficient active algorithms for hierarchical clustering. *arXiv preprint arXiv:1206.4672*, 2012.
- [MB06] N. Meinshausen and P. Bühlmann. High dimensional graphs and variable selection with the lasso. *Annals of Statistics*, 34(3):1436–1462, 2006.
- [Mic12] Michael. Boruvka algorithm parallel implementation cuda, December 2012.
- [Mos07] Elchanan Mossel. Distorted metrics on trees and phylogenetic forests. *IEEE/ACM Transactions on Computational Biology and Bioinformatics (TCBB)*, 4(1):108–116, 2007.
- [MR05] Elchanan Mossel and Sébastien Roch. Learning nonsingular phylogenies and hidden markov models. In *Proceedings of the thirty-seventh annual ACM symposium on Theory of computing*, pages 366–375. ACM, 2005.
- [Pea88] Judea Pearl. *Probabilistic reasoning in intelligent systems: networks of plausible inference*. Morgan Kaufmann, 1988.

- [RF81] David F. Robinson and Leslie R. Foulds. Comparison of phylogenetic trees. *Mathematical Biosciences*, 53(1):131–147, 1981.
- [SN10] Alexander Smola and Shravan Narayanamurthy. An architecture for parallel topic models. *Proceedings of the VLDB Endowment*, 3(1-2):703–710, 2010.
- [VHPN09] Vibhav Vineet, Pawan Harish, Suryakant Patidar, and P. J. Narayanan. Fast minimum spanning tree for large graphs on the gpu. In *Proceedings of the Conference on High Performance Graphics 2009*, pages 167–171. ACM, 2009.
- [WDK⁺13] J. Wei, W. Dai, A. Kumar, X. Zheng, Q. Ho, and E. P. Xing. Consistent Bounded-Asynchronous Parameter Servers for Distributed ML. *ArXiv e-prints*, December 2013.
- [WL13] Fang Wang and Yi Li. Beyond physical connections: Tree models in human pose estimation. In *Proc. of CVPR*, 2013.

Appendix

A Additivity of the Multivariate Information Distance

Recall that the additive information distance between nodes two categorical variables x_i and x_j was defined in [CTAW11]. We extend the notation of information distance to high dimensional variables via Definition 4.1 and present the proof of its additivity in Lemma 4.2 here.

Proof.

$$\mathbb{E}[x_a x_c^\top] = \mathbb{E}[\mathbb{E}[x_a x_c^\top | x_b]] = A \mathbb{E}[x_b x_b^\top] B^\top$$

Consider three nodes a, b, c such that there are edges between a and b , and b and c . Let the $A = \mathbb{E}(x_a | x_b)$ and $B = \mathbb{E}(x_c | x_b)$. From Definition 4.1, we have, assuming that $\mathbb{E}(x_a x_a^\top)$, $\mathbb{E}(x_b x_b^\top)$ and $\mathbb{E}(x_c x_c^\top)$ are full rank.

$$\begin{aligned} \text{dist}(v_a, v_c) &= -\log \frac{\prod_{i=1}^k \sigma_i(\mathbb{E}(x_a x_c^\top))}{\sqrt{\det(\mathbb{E}(x_a x_a^\top)) \det(\mathbb{E}(x_c x_c^\top))}} \\ e^{-\text{dist}(v_a, v_c)} &= \det \left(\mathbb{E}(x_a x_a^\top)^{-1/2} U^\top \mathbb{E}(x_a x_c^\top) V \mathbb{E}(x_c x_c^\top)^{-1/2} \right) \end{aligned}$$

where k -SVD($\mathbb{E}(x_a x_c^\top)$) = $U \Sigma V^\top$. Similarly,

$$\begin{aligned} e^{-\text{dist}(v_a, v_b)} &= \det \left(\mathbb{E}(x_a x_a^\top)^{-1/2} U^\top \mathbb{E}(x_a x_b^\top) W \mathbb{E}(x_b x_b^\top)^{-1/2} \right) \\ e^{-\text{dist}(v_b, v_c)} &= \det \left(\mathbb{E}(x_b x_b^\top)^{-1/2} W^\top \mathbb{E}(x_b x_c^\top) V \mathbb{E}(x_c x_c^\top)^{-1/2} \right) \end{aligned}$$

where k -SVD($\mathbb{E}(x_a x_b^\top)$) = $U \Sigma W^\top$ and k -SVD($\mathbb{E}(x_b x_c^\top)$) = $W \Sigma V^\top$.

Therefore,

$$\begin{aligned} e^{-(\text{dist}(a,b) + \text{dist}(b,c))} &= \det(\mathbb{E}(x_a x_a^\top)^{-1/2} U^\top \mathbb{E}(x_a x_b^\top) \mathbb{E}(x_b x_b^\top)^{-1/2-1/2} \mathbb{E}(x_b x_c^\top) V \mathbb{E}(x_c x_c^\top)^{-1/2}) \\ &= \det(\mathbb{E}(x_a x_a^\top)^{-1/2} U^\top A \mathbb{E}(x_b x_b^\top) B^\top V \mathbb{E}(x_c x_c^\top)^{-1/2}) = e^{-\text{dist}(v_a, v_c)} \end{aligned}$$

We conclude that the multivariate information distance is additive. Note that $\mathbb{E}[x_a x_b^\top] = \mathbb{E}(\mathbb{E}(x_a x_b^\top | x_b)) = \mathbb{E}(A x_b x_b^\top) = A \mathbb{E}(x_b x_b^\top)$. \square

We note that when the second moments are not full rank, the above distance can be extended as follows:

$$\text{dist}(v_a, v_c) = -\log \frac{\prod_{i=1}^k \sigma_i(\mathbb{E}(x_a x_c^\top))}{\sqrt{\prod_{i=1}^k \sigma_i(\mathbb{E}(x_a x_a^\top)) \prod_{i=1}^k \sigma_i(\mathbb{E}(x_c x_c^\top))}}.$$

B Local Recursive Grouping

The Local Recursive Grouping (LRG) algorithm is a local divide and conquer procedure for learning the structure and parameter of the latent tree (Algorithm 1). We perform recursive grouping simultaneously on the sub-trees of the MST. Each of the sub-tree consists of an internal node and its neighborhood nodes. We keep track of the internal nodes of the MST, and their neighbors. The resultant latent sub-trees after LRG can be merged easily to recover the final latent tree. Consider a pair of neighboring sub-trees in the MST. They have two common nodes (the internal nodes) which are neighbors on MST. Firstly we identify the path from one internal node to the other in the trees to be merged, then compute the multivariate information distances between the internal nodes and the introduced hidden

nodes. We recover the path between the two internal nodes in the merged tree by inserting the hidden nodes closely to their surrogate node. Secondly, we merge all the leaves which are not in this path by attaching them to their parent. Hence, the recursive grouping can be done in parallel and we can recover the latent tree structure via this merging method.

Lemma B.1. *If an observable node v_j is the surrogate node of a hidden node h_i , then the hidden node h_i can be discovered using v_j and the neighbors of v_j in the MST.*

This is due to the additive property of the multivariate information distance on the tree and the definition of a surrogate node. This observation is crucial for a completely local and parallel structure and parameter estimation. It is also easy to see that all internal nodes in the MST are surrogate nodes.

After the parallel construction of the MST, we look at all the internal nodes \mathcal{X}_{int} . For $v_i \in \mathcal{X}_{\text{int}}$, we denote the neighborhood of v_i on MST as $\text{nbd}_{\text{sub}}(v_i; \text{MST})$ which is a small sub-tree. Note that the number of such sub-trees is equal to the number of internal nodes in MST.

For any pair of sub-trees, $\text{nbd}_{\text{sub}}(v_i; \text{MST})$ and $\text{nbd}_{\text{sub}}(v_j; \text{MST})$, there are two topological relationships, namely overlapping (i.e., when the sub-trees share at least one node in common) and non-overlapping (i.e., when the sub-trees do not share any nodes).

Since we define a neighborhood centered at v_i as only its immediate neighbors and itself on MST, the overlapping neighborhood pair $\text{nbd}_{\text{sub}}(v_i; \text{MST})$ and $\text{nbd}_{\text{sub}}(v_j; \text{MST})$ can only have conflicting paths, namely $\text{path}(v_i, v_j; \mathcal{N}_i)$ and $\text{path}(v_i, v_j; \mathcal{N}_j)$, if v_i and v_j are neighbors in MST.

With this in mind, we locally estimate all the latent sub-trees, denoted as \mathcal{N}_i , by applying Recursive Grouping [CTAW11] in a parallel manner on $\text{nbd}_{\text{sub}}(v_i; \text{MST})$, $\forall v_i \in \mathcal{X}_{\text{int}}$. Note that the latent nodes automatically introduced by $\text{RG}(v_i)$ have v_i as their surrogate. We update the tree structure by joining each level in a bottom-up manner. The testing of the relationship among nodes [CTAW11] uses the additive multivariate information distance metric (Appendix A) $\Phi(v_i, v_j; k) = \text{dist}(v_i, v_k) - \text{dist}(v_i, v_k)$ to decide whether the nodes v_i and v_j are parent-child or siblings. If they are siblings, they should be joined by a hidden parent. If they are parent and child, the child node is placed as a lower level node and we add the other node as the single parent node, which is then joined in the next level.

Finally, for each internal edge of MST connecting two internal nodes v_i and v_j , we consider merging the latent sub-trees. In the example of two local estimated latent sub-trees in Figure 1, we illustrate the complete local merging algorithm that we propose.

C Proof Sketch for Theorem 7.1

We argue for the correctness of the method under exact moments. The sample complexity follows from the previous works. In order to clarify the proof ideas, we define the notion of *surrogate node* [CTAW11] as follows.

Definition C.1. *Surrogate node for hidden node h_i on the latent tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ is defined as $\text{Sg}(h_i; \mathcal{T}) := \arg \min_{v_j \in \mathcal{X}} \text{dist}(v_i, v_j)$.*

In other words, the surrogate for a hidden node is an observable node which has the minimum multivariate information distance from the hidden node. See Figure 1(a), the surrogate node of h_1 , $\text{Sg}(h_1; \mathcal{T})$, is v_3 , $\text{Sg}(h_2; \mathcal{T}) = \text{Sg}(h_3; \mathcal{T}) = v_5$. Note that the notion of the surrogate node is only required for analysis, and our algorithm does not need to know this information.

The notion of surrogacy allows us to relate the constructed MST (over observed nodes) with the underlying latent tree. It can be easily shown that contracting the hidden nodes to their surrogates on latent tree leads to MST. Local recursive grouping procedure can be viewed as reversing these contractions, and hence, we obtain consistent local sub-trees.

We now argue the correctness of the structure union procedure, which merges the local sub-trees. In each reconstructed sub-tree \mathcal{N}_i , where v_i is the group leader, the discovered hidden nodes $\{h^i\}$ form a surrogate relationship with v_i , i.e. $\text{Sg}(h^i; \mathcal{T}) = v_i$. Our merging approach maintains these surrogate relationships. For example in Figure 1(d1,d2), we have the path $v_3 - h_1 - v_5$ in \mathcal{N}_3 and path $v_3 - h_3 - h_2 - v_5$ in \mathcal{N}_5 . The resulting path is

$v_3 - h_1 - h_3 - h_2 - v_5$, as seen in Figure 1(e). We now argue why this is correct. As discussed before, $\text{Sg}(h_1; \mathcal{T}) = v_3$ and $\text{Sg}(h_2; \mathcal{T}) = \text{Sg}(h_3; \mathcal{T}) = v_5$. When we merge the two subtrees, we want to preserve the paths from the group leaders to the added hidden nodes, and this ensures that the surrogate relationships are preserved in the resulting merged tree. Thus, we obtain a global consistent tree structure by merging the local structures. The correctness of parameter learning comes from the consistency of the tensor decomposition techniques and careful alignments of the hidden labels across different decompositions. Refer to Appendix D, G for proof details and the sample complexity.

D Proof of Correctness for LRG

Definition D.1. A latent tree $\mathcal{T}_{\geq 3}$ is defined to be a minimal (or identifiable) latent tree if it satisfies that each latent variable has at least 3 neighbors.

Definition D.2. Surrogate node for hidden node h_i in latent tree $\mathcal{T} = (\mathcal{V}, \mathcal{E})$ is defined as

$$\text{Sg}(h_i; \mathcal{T}) := \arg \min_{v_j \in \mathcal{X}} \text{dist}(v_i, v_j).$$

There are some useful observations about the MST in [CTAW11] which we recall here.

Property D.3 (MST – surrogate neighborhood preservation). *The surrogate nodes of any two neighboring nodes in \mathcal{E} are also neighbors in the MST. I.e.,*

$$(h_i, h_j) \in \mathcal{E} \Rightarrow (\text{Sg}(h_i), \text{Sg}(h_j)) \in \text{MST}.$$

Property D.4 (MST – surrogate consistency along path). *If $v_j \in \mathcal{X}$ and $v_h \in \text{Sg}^{-1}(v_j)$, then every node along the path connecting v_j and v_h belongs to the inverse surrogate set $\text{Sg}^{-1}(v_j)$, i.e.,*

$$v_i \in \text{Sg}^{-1}(v_j), \forall v_i \in \text{Path}(v_j, v_h)$$

if

$$v_h \in \text{Sg}^{-1}(v_j).$$

The MST properties observed connect the MST over observable nodes with the original latent tree \mathcal{T} . We obtain MST by contracting all the latent nodes to its surrogate node.

Given that the correctness of CLR algorithm is proved in [CTAW11], we prove the equivalence between the CLR and PLR.

Lemma D.5. *For any sub-tree pairs $\text{nb}[v_i; \text{MST}]$ and $\text{nb}[v_j; \text{MST}]$, there is at most one overlapping edge. The overlapping edge exists if and only if $v_i \in \text{nb}(v_j; \text{MST})$.*

This is easy to see.

Lemma D.6. *Denote the latent tree recovered from $\text{nb}[v_i; \text{MST}]$ as \mathcal{N}_i and similarly for $\text{nb}[v_j; \text{MST}]$. The inconsistency, if any, between \mathcal{N}_i and \mathcal{N}_j occurs in the overlapping path $(v_i, v_j; \mathcal{N}_i)$ in and path $(v_i, v_j; \mathcal{N}_j)$ after LRG implementation on each subtrees.*

We now prove the correctness of LRG. Let us denote the latent tree resulting from merging a subset of small latent trees as $T_{\text{LRG}}(S)$, where S is the set of center of subtrees that are merged pair-wisely. CLR algorithm in [CTAW11] implements the RG in a serial manner. Let us denote the latent tree learned at iteration i from CLR is $T_{\text{CLR}}(S)$, where S is the set of internal nodes visited by CLR at current iteration. We prove the correctness of LRG by induction on the iterations.

At the initial step $S = \emptyset$: $T_{\text{CLR}} = \text{MST}$ and $T_{\text{LRG}} = \text{MST}$, thus $T_{\text{CLR}} = T_{\text{LRG}}$.

Now we assume that for the same set S_{i-1} , $T_{\text{CLR}} = T_{\text{LRG}}$ is true for $r = 1, \dots, i-1$. At iteration $r = i$ where CLR employs RG on the immediate neighborhood of node v_i on $T_{\text{CLR}}(S_{i-1})$, let us assume that H_i is the set of hidden nodes who are immediate neighbors of $i-1$. The CLR algorithm thus considers all the neighbors and

implements the RG. We know that the surrogate nodes of every latent node in H_i belong to previously visited nodes S_{i-1} . According to Property D.3 and D.4, if we contract all the hidden node neighbors to their surrogate nodes, CLRG thus is a RG on neighborhood of i on MST.

As for our LRG algorithm at this step, $T_{\text{LRG}}(S_i)$ is the merging between $T_{\text{LRG}}(S_{i-1})$ and \mathcal{N}_i . The latent nodes whose surrogate node is j are introduced between the edge $(i-1, i)$. Now that we know \mathcal{N}_i is the RG output from immediate neighborhood of i on MST. Therefore, we proved that $T_{\text{CLRG}}(S_i) = T_{\text{LRG}}(S_i)$.

E Cross Group Alignment Correction

In order to achieve cross group alignments, tensor decompositions on two cross group triplets have to be computed. The first triplet is formed by three nodes: reference node in group 1, x_1 , non-reference node in group 1, x_2 , and reference node in group 2, x_3 . The second triplet is formed by three nodes as well: reference node in group 2, x_3 , non-reference node in group 2, x_4 and reference node in group 1, x_1 . Let us use h_1 to denote the parent node in group 1, and h_2 the parent node in group 2.

From $\text{Trip}(x_1, x_2, x_3)$, we obtain $P(h_1|x_1) = \tilde{A}$, $P(x_2|h_1) = B$ and $P(x_3|h_1) = P(x_3|h_2)P(h_2|h_1) = DE$. From $\text{Trip}(x_3, x_4, x_1)$, we know $P(x_3|h_2) = D\Pi$, $P(x_4|h_2) = C\Pi$ and $P(h_2|x_1) = P(h_2|h_1)P(h_1|x_1) = \Pi E \tilde{A}$, where Π is a permutation matrix. We compute Π as $\Pi = \sqrt{(\Pi E \tilde{A})(\tilde{A})^\dagger (DE)^\dagger (D\Pi)}$ so that $D = (D\Pi)\Pi^\dagger$ is aligned with group 1. Thus, when all the parameters in the two groups are aligned by permute group 2 parameters using Π , thus the alignment is completed.

Similarly, the alignment correction can be done by calculating the permutation matrices while merging different threads.

Overall, we merge the local structures and align the parameters from LRG local sub-trees using Procedure 2 and 3.

F Computational Complexity

We recall some notations here: d is the observable node dimension, k is the hidden node dimension ($k \ll d$), N is the number of samples, p is the number of observable nodes, and z is the number of non-zero elements in each sample.

Multivariate information distance estimation involves sparse matrix multiplications to compute the pairwise second moments. Each observable node has a $d \times N$ sample matrix with z non-zeros per column. Computing the product $x_1 x_2^T$ from a single sample for nodes 1 and 2 requires $O(z)$ time and there are N such sample pair products leading to $O(Nz)$ time. There are $O(p^2)$ node pairs and hence the degree of parallelism is $O(p^2)$. Next, we perform the k -rank SVD of each of these matrices. Each SVD takes $O(d^2 k)$ time using classical methods. Using randomized methods [GM13a], this can be improved to $O(d + k^3)$.

Next on, we construct the MST in $O(\log p)$ time per worker with p^2 workers. The structure learning can be done in $O(\Gamma^3)$ per sub-tree and the local neighborhood of each node can be processed completely in parallel. We assume that the group sizes Γ are constant (the sizes are determined by the degree of nodes in the latent tree and homogeneity of parameters across different edges of the tree). The parameter estimation of each triplet of nodes consists of implicit stochastic updates involving products of $k \times k$ and $d \times k$ matrices. Note that we do not need to consider all possible triplets in groups but each node must be take care by a triplet and hence there are $O(p)$ triplets. This leads to a factor of $O(\Gamma k^3 + \Gamma d k^2)$ time per worker with p/Γ degree of parallelism.

At last, the merging step consists of products of $k \times k$ and $d \times k$ matrices for each edge in the latent tree leading to $O(d k^2)$ time per worker with p/Γ degree of parallelism.

G Sample Complexity

From [ACH⁺11], we recall the number of samples required for the recovery of the tree structure that is consistent with the ground truth (for a precise definition of consistency, refer to Definition 2 of [CTAW11]).

Lemma G.1. *If*

$$N > \frac{200k^2 B^2 t}{\left(\frac{\gamma_{\min}^2}{\gamma_{\max}}(1 - \text{dist}_{\max})\right)^2} + \frac{7kM^2 t}{\frac{\gamma_{\min}^2}{\gamma_{\max}}(1 - \text{dist}_{\max})}, \quad (9)$$

then with probability at least $1 - \eta$, proposed algorithm returns $\hat{\mathcal{T}} = \mathcal{T}$, where

$$\begin{aligned} B &:= \max_{x_i, x_j \in \mathcal{X}} \left\{ \sqrt{\max\{\|\mathbb{E}[\|x_i\|^2 x_j x_j^\top]\|, \max\{\|\mathbb{E}[\|x_j\|^2 x_i x_i^\top]\|\}} \right\}, \\ M &:= \max_{x_i \in \mathcal{X}} \{\|x_i\|\}, \\ t &:= \max_{x_i, x_j \in \mathcal{X}} \left\{ 4 \ln(4) \frac{\mathbb{E}[\|x_i\|^2 \|x_j\|^2] - \text{Tr}(\mathbb{E}[x_i x_i^\top] \mathbb{E}[x_j x_j^\top])}{\max\{\|\mathbb{E}[\|x_j\|^2 x_i x_i^\top]\|, \|\mathbb{E}[\|x_i\|^2 x_j x_j^\top]\|\}} n/\eta \right\}. \\ \gamma_{\min} &:= \min_{\{x_1, x_2\}} \{\sigma(\mathbb{E}[x_1 x_2^\top])\} \\ \gamma_{\max} &:= \max_{\{x_1, x_2\}} \{\sigma(\mathbb{E}[x_1 x_2^\top])\} \end{aligned}$$

From [AFH⁺12], we recall the sample complexity for the faithful recovery of parameters via tensor decomposition methods.

We define ϵ_P to be the noise raised between empirical estimation of the second order moments and exact second order moments, and ϵ_T to be the noise raised between empirical estimation of the third order moments and the exact third order moments.

Lemma G.2. *Consider positive constants C, C', c and c' , the following holds. If*

$$\begin{aligned} \epsilon_P &\leq c \frac{\lambda_k}{k}, \quad \epsilon_T \leq c' \frac{\lambda_k \sigma_k^{3/2}}{k} \\ N &\geq C \left(\log(k) + \log \left(\log \left(\frac{\lambda_1 \sigma_k^{3/2}}{\epsilon_T} + \frac{1}{\epsilon_P} \right) \right) \right) \\ L &\geq \text{poly}(k) \log(1/\delta), \end{aligned}$$

then with probability at least $1 - \delta$, tensor decomposition returns $(\hat{v}_i, \hat{\lambda}_i) : i \in [k]$ satisfying, after appropriate reordering,

$$\begin{aligned} \|\hat{v}_i - v_i\|_2 &\leq C' \left(\frac{1}{\lambda_i} \frac{1}{\sigma_k^2} \epsilon_T + \left(\frac{\lambda_1}{\lambda_i} \frac{1}{\sqrt{\sigma_k}} + 1 \right) \epsilon_P \right) \\ |\hat{\lambda}_i - \lambda_i| &\leq C' \left(\frac{1}{\sigma_k^{3/2}} \epsilon_T + \lambda_1 \epsilon_P \right) \end{aligned}$$

for all $i \in [k]$.

We note that $\sigma_1 \geq \sigma_2 \geq \dots \sigma_k > 0$ are the non-zero singular values of the second order moments, $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_k > 0$ are the ground-truth eigenvalues of the third order moments, and v_i are the corresponding eigenvectors for all $i \in [k]$.

H Efficient SVD Using Sparsity and Dimensionality Reduction

Without loss of generality, we assume that a matrix whose SVD we aim to compute has no row or column which is fully zeros, since, if it does have zero entries, such row and columns can be dropped.

Let $A \in \mathbb{R}^{n \times n}$ be the matrix to do SVD. Let $\Phi \in \mathbb{R}^{d \times \tilde{k}}$, where $\tilde{k} = \alpha k$ with α is a scalar, usually, in the range $[2, 3]$. For the i^{th} row of Φ , if $\sum_i |\Phi|(i, :) \neq 0$ and $\sum_i |\Phi|(:, i) \neq 0$, then there is only one non-zero entry and that entry is uniformly chosen from $[\tilde{k}]$. If either $\sum_i |\Phi|(i, :) = 0$ or $\sum_i |\Phi|(:, i) = 0$, we leave that row blank. Let $D \in \mathbb{R}^{d \times d}$ be a diagonal matrix with iid Rademacher entries, i.e., each non-zero entry is 1 or -1 with probability $\frac{1}{2}$. Now, our embedding matrix [CW13] is $S = D\Phi$, i.e., we find AS and then proceed with the Nystrom [HNN⁺13] method. Unlike the usual Nystrom method [GM13b] which uses a random matrix for computing the embedding, we improve upon this by using a sparse matrix for the embedding since the sparsity improves the running time and the memory requirements of the algorithm.